



Replication

ONTAP 9

NetApp
February 12, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap/concepts/snapshot-copies-concept.html> on February 12, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Replication	1
S snapshots	1
SnapMirror disaster recovery and data transfer	2
SnapMirror cloud backups to object storage	3
SnapVault archiving	4
Cloud backup and support for traditional backups	5
MetroCluster continuous availability	6

Replication

Snapshots

Traditionally, ONTAP replication technologies served the need for disaster recovery (DR) and data archiving. With the advent of cloud services, ONTAP replication has been adapted to data transfer between endpoints in the NetApp data fabric. The foundation for all these uses is ONTAP snapshot technology.

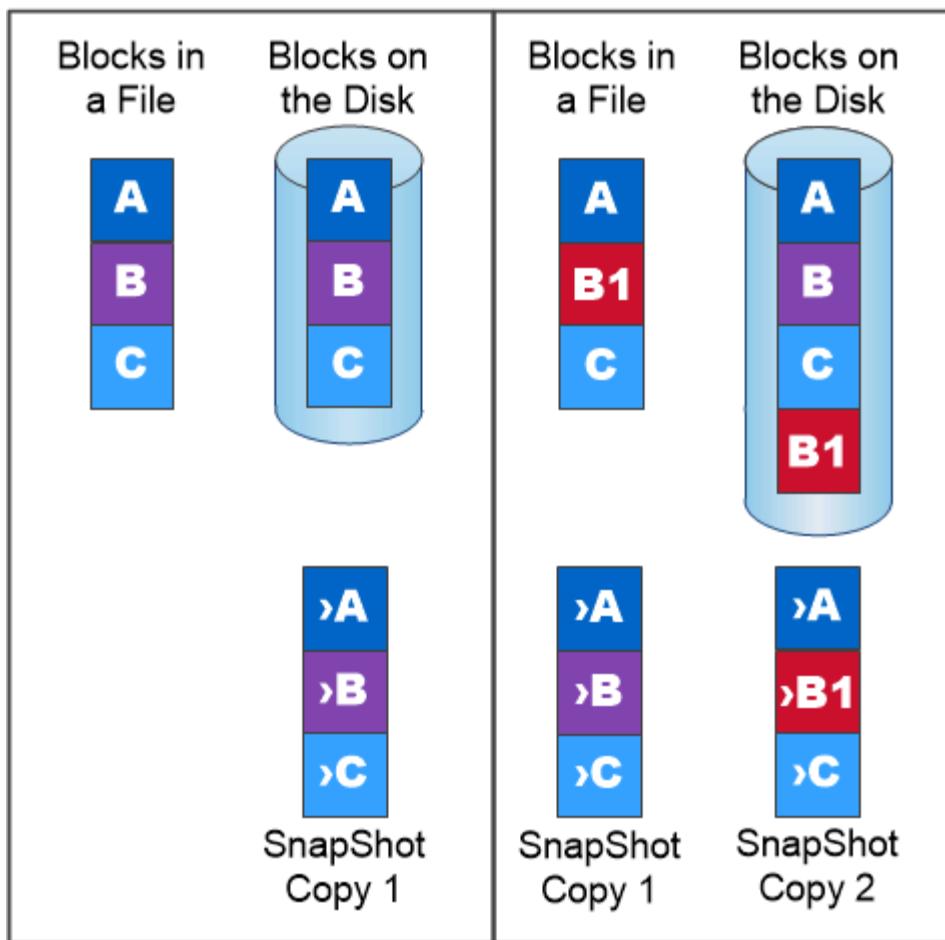
A *snapshot* (formerly *Snapshot copy*) is a read-only, point-in-time image of a volume. After a snapshot is created, the active file system and the snapshot point to the same disk blocks; therefore, the snapshot doesn't use extra disk space. Over time, the image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last snapshot was made.

Snapshots owe their efficiency to ONTAP's core storage virtualization technology, its *Write Anywhere File Layout (WAFL)*. Like a database, WAFL uses metadata to point to actual data blocks on disk. But, unlike a database, WAFL does not overwrite existing blocks. It writes updated data to a new block and changes the metadata.

Snapshots are efficient because, rather copy data blocks, ONTAP references metadata when creating a snapshot. Doing so eliminates both the "seek time" that other systems incur in locating the blocks to copy and the cost of making the copy itself.

You can use a snapshot to recover individual files or LUNs, or to restore the entire contents of a volume. ONTAP compares pointer information in the snapshot with data on disk to reconstruct the missing or damaged object, without downtime or a significant performance cost.

A *snapshot policy* defines how the system creates snapshots of volumes. The policy specifies when to create the snapshots, how many copies to retain, how to name them, and how to label them for replication. For example, a system might create one snapshot every day at 12:10 a.m., retain the two most recent copies, name them "daily" (appended with a timestamp), and label them "daily" for replication.



A *Snapshot copy* records only changes to the active file system since the last *Snapshot copy*.

SnapMirror disaster recovery and data transfer

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. As its name implies, *SnapMirror* creates a replica, or *mirror*, of your working data in secondary storage from which you can continue to serve data in the event of a catastrophe at the primary site.

Data is mirrored at the volume level. The relationship between the source volume in primary storage and the destination volume in secondary storage is called a *data protection relationship*. The clusters in which the volumes reside and the SVMs that serve data from the volumes must be *peered*. A peer relationship enables clusters and SVMs to exchange data securely.



You can also create a data protection relationship between SVMs. In this type of relationship, all or part of the SVM's configuration, from NFS exports and SMB shares to RBAC, is replicated, as well as the data in the volumes the SVM owns.

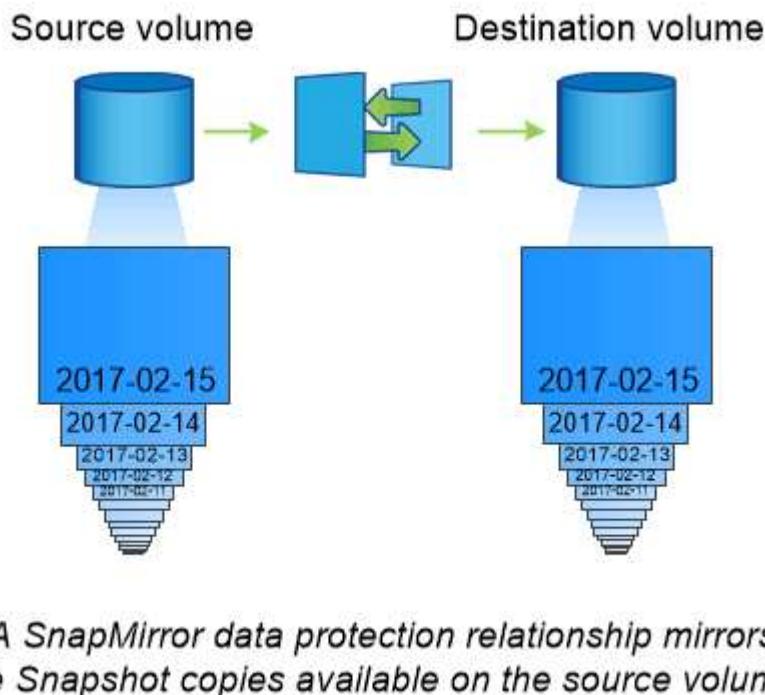
Beginning with ONTAP 9.10.1, you can create data protection relationships between S3 buckets using SnapMirror S3. Destination buckets can be on local or remote ONTAP systems, or on non-ONTAP systems such as StorageGRID and AWS.

The first time you invoke SnapMirror, it performs a *baseline transfer* from the source volume to the destination volume. The baseline transfer typically involves the following steps:

- Make a snapshot of the source volume.
- Transfer the snapshot and all the data blocks it references to the destination volume.
- Transfer the remaining, less recent snapshots on the source volume to the destination volume for use in case the “active” mirror is corrupted.

Once a baseline transfer is complete, SnapMirror transfers only new snapshots to the mirror. Updates are asynchronous, following the schedule you configure. Retention mirrors the snapshot policy on the source. You can activate the destination volume with minimal disruption in case of a disaster at the primary site, and reactivate the source volume when service is restored.

Because SnapMirror transfers only snapshots after the baseline is created, replication is fast and nondisruptive. As the failover use case implies, the controllers on the secondary system should be equivalent or nearly equivalent to the controllers on the primary system to serve data efficiently from mirrored storage.



Using SnapMirror for data transfer

You can also use SnapMirror to replicate data between endpoints in the NetApp data fabric. You can choose between one-time replication or recurring replication when you create the SnapMirror policy.

SnapMirror cloud backups to object storage

SnapMirror cloud is a backup and recovery technology designed for ONTAP users who want to transition their data protection workflows to the cloud. Organizations moving away from legacy backup-to-tape architectures can use object storage as an alternative repository for long-term data retention and archiving. SnapMirror cloud provides ONTAP-to-object storage replication as part of an incremental forever backup strategy.

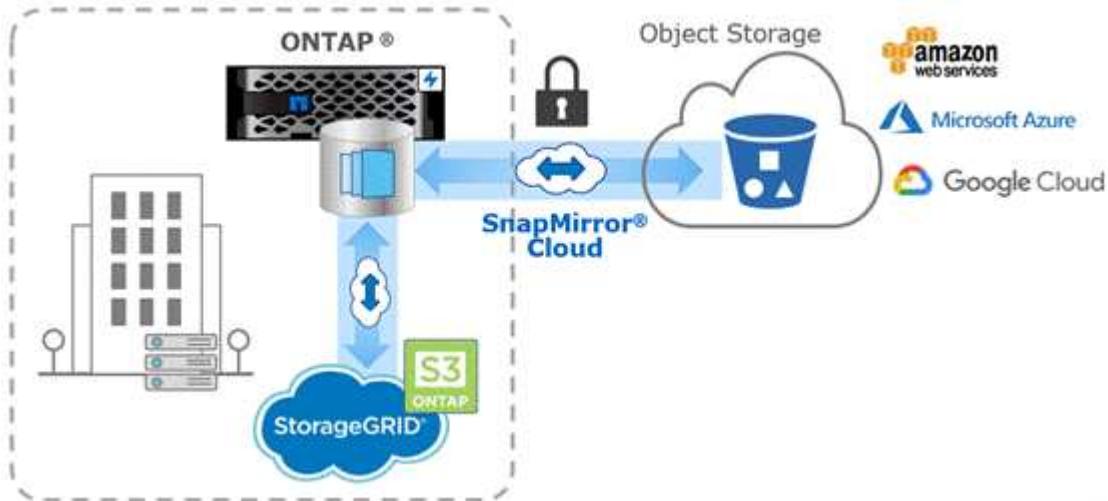
SnapMirror cloud replication is a licensed ONTAP feature. SnapMirror cloud was introduced in ONTAP 9.8 as an extension to the family of SnapMirror replication technologies. While SnapMirror is frequently used for ONTAP-to-ONTAP backups, SnapMirror cloud uses the same replication engine to transfer snapshots for ONTAP to S3-compliant object storage backups.

Targeted for backup use cases, SnapMirror cloud supports both long-term retention and archives workflows. As with SnapMirror, the initial SnapMirror cloud backup performs a baseline transfer of a volume. For subsequent backups, SnapMirror cloud generates a snapshot of the source volume and transfers the snapshot with only the changed data blocks to an object storage target.

SnapMirror cloud relationships can be configured between ONTAP systems and select on-premises and public cloud object storage targets - including Amazon S3, Google Cloud Storage, and Microsoft Azure Blob Storage. Additional on-premises object storage targets include StorageGRID and ONTAP S3.

In addition to using ONTAP System Manager to manage SnapMirror cloud configurations, several orchestration options are available for managing SnapMirror cloud backups:

- Multiple 3rd party backup partners who offer support for SnapMirror cloud replication. Participating vendors are available on the [NetApp blog](#).
- NetApp Backup and Recovery for a NetApp-native solution for ONTAP environments
- APIs for developing custom software for data protection workflows or leveraging automation tools



SnapVault archiving

The SnapMirror license is used to support both SnapVault relationships for backup, and SnapMirror relationships for disaster recovery. Beginning with ONTAP 9.3, SnapVault licenses are deprecated, and SnapMirror licenses can be used to configure vault, mirror, and mirror-and-vault relationships. SnapMirror replication is used for ONTAP-to-ONTAP replication of snapshots, supporting both backup and disaster recovery use cases.

SnapVault is archiving technology, designed for disk-to-disk snapshot replication for standards compliance and other governance-related purposes. In contrast to a SnapMirror relationship, in which the destination usually contains only the snapshots currently in the source volume, a SnapVault destination typically retains point-in-time snapshots created over a much longer period.

You might want to keep monthly snapshots of your data over a 20-year span, for example, to comply with

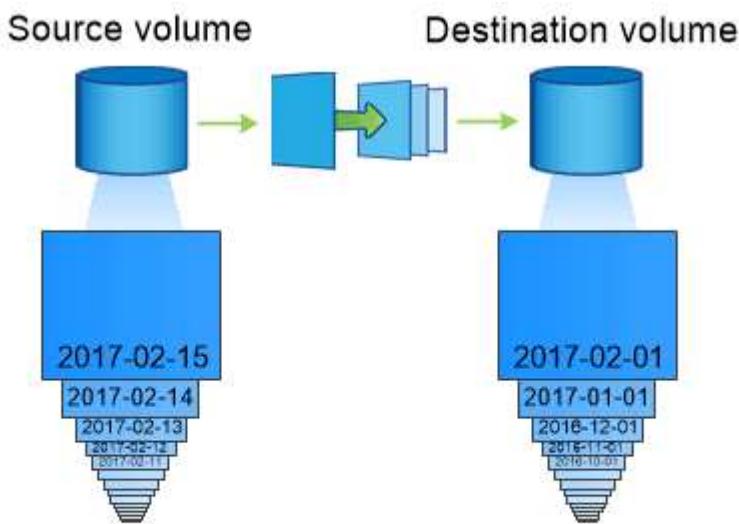
government accounting regulations for your business. Since there is no requirement to serve data from vault storage, you can use slower, less expensive disks on the destination system.

As with SnapMirror, SnapVault performs a baseline transfer the first time you invoke it. It makes a snapshot of the source volume, then transfers the copy and the data blocks it references to the destination volume. Unlike SnapMirror, SnapVault does not include older snapshots in the baseline.

Updates are asynchronous, following the schedule you configure. The rules you define in the policy for the relationship identify which new snapshots to include in updates and how many copies to retain. The labels defined in the policy (“monthly,” for example) must match one or more labels defined in the snapshot policy on the source. Otherwise, replication fails.



SnapMirror and SnapVault share the same command infrastructure. You specify which method you want to use when you create a policy. Both methods require peered clusters and peered SVMs.



A SnapVault data protection relationship typically retains point-in-time Snapshot copies created over a longer period than the Snapshot copies on the source volume.

Cloud backup and support for traditional backups

In addition to SnapMirror and SnapVault data protection relationships, which were disk-to-disk only for ONTAP 9.7 and earlier, there are now several backup solutions that offer a less expensive alternative for long-term data retention.

Numerous third-party data protection applications offer traditional backup for ONTAP-managed data. Veeam, Veritas, and Commvault, among others, all offer integrated backup for ONTAP systems.

Beginning with ONTAP 9.8, SnapMirror cloud provides asynchronous replication of snapshots from ONTAP instances to object storage endpoints. SnapMirror cloud replication requires a licensed application for orchestration and management of data protection workflows. SnapMirror cloud relationships are supported from ONTAP systems to select on-premises and public cloud object storage targets — including AWS S3, Google Cloud Storage Platform, or Microsoft Azure Blob Storage — which provides enhanced efficiency with vendor backup software. Contact your NetApp representative for a list of supported certified applications and

object storage vendors.

If you are interested in cloud-native data protection, the NetApp Console can be used to configure SnapMirror or SnapVault relationships between on-premises volumes and Cloud Volumes ONTAP instances in the public cloud.

The Console also provides backups of Cloud Volumes ONTAP instances using a Software as a Service (SaaS) model. Users can back up their Cloud Volumes ONTAP instances to S3 and S3-compliant public cloud object storage using NetApp Backup and Recovery.

[Cloud Volumes ONTAP documentation](#)

[NetApp Console documentation](#)

[NetApp Console](#)

MetroCluster continuous availability

MetroCluster configurations protect data by implementing two physically separate, mirrored clusters. Each cluster synchronously replicates the data and SVM configuration of the other. In the event of a disaster at one site, an administrator can activate the mirrored SVM and begin serving data from the surviving site.

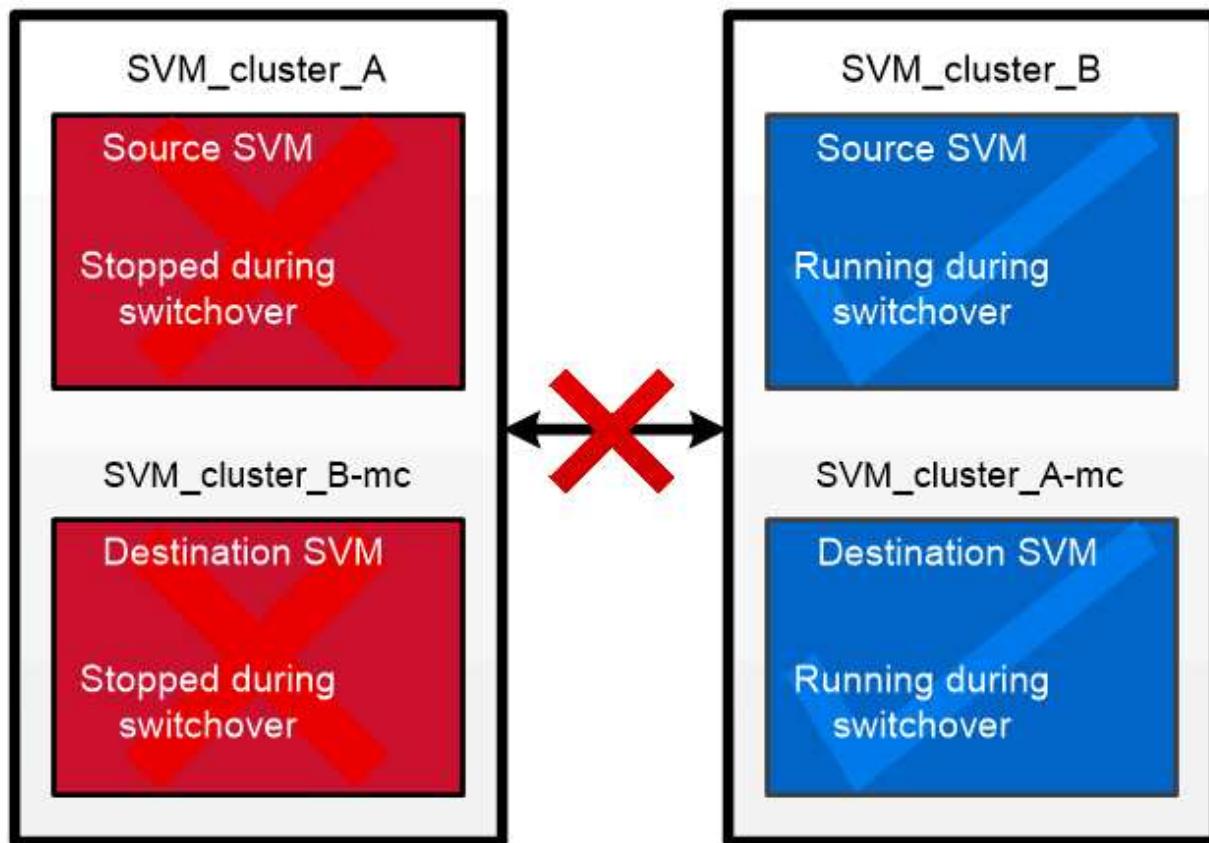
- *Fabric-attached MetroCluster* and *MetroCluster IP* configurations support metropolitan-wide clusters.
- *Stretch MetroCluster* configurations support campus-wide clusters.

Clusters must be peered in either case.

MetroCluster uses an ONTAP feature called *SyncMirror* to synchronously mirror aggregate data for each cluster in copies, or *plexes*, in the other cluster's storage. If a switchover occurs, the remote plex on the surviving cluster comes online and the secondary SVM begins serving data.

Cluster A down and
switched over

Cluster B up



When a MetroCluster switchover occurs, the remote plex on the surviving cluster comes online and the secondary SVM begins serving data.

Using SyncMirror in non-MetroCluster implementations

You can optionally use SyncMirror in a non-MetroCluster implementation to protect against data loss if more disks fail than the RAID type protects against, or if there is a loss of connectivity to RAID group disks. The feature is available for HA pairs only.

Aggregate data is mirrored in plexes stored on different disk shelves. If one of the shelves becomes unavailable, the unaffected plex continues to serve data while you fix the cause of the failure.

Keep in mind that an aggregate mirrored using SyncMirror requires twice as much storage as an unmirrored aggregate. Each plex requires as many disks as the plex it mirrors. You would need 2,880 GB of disk space, for example, to mirror a 1,440 GB aggregate, 1,440 GB for each plex.

With SyncMirror, it's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space may be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write snapshot-based architecture. Failure to adhere to these best practices may have a negative impact on SyncMirror resynchronization performance, which indirectly impacts operational workflows such as NDU for non-shared cloud deployments and switchback for MetroCluster deployments.

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.